

STATISTICS WORKSHOP III

United States Department of Agriculture

Hypothesis Tests and Interval Estimation

presented by
Vicki Lancaster
NEPTUNE AND CO., INC.
vlancast@neptuneandco.com

Interval Estimation & Hypothesis

Goals


- " Introduce the terminology of *Type I error, Type II error, power, test statistic, significance level, and margin of error.*
- " Conduct *hypothesis test* intervals for comparing a sample for hypothesized value.
- " Explore the relationship between *hypothesis test* and a *confidence interval*.
- " Explore the relationship between a *one-sided hypothesis test* and a *confidence limit*.

Interval Estimation & Hypothesis Testing

Statistical inference is a formal process of drawing conclusions from sample data that take into account the effects of chance variation. Probability is used to quantify how confident the researcher is that conclusions drawn from the sample are not the result of a chance occurrence.

The objective of an inferential method is to draw inferences about a population from information contained in a sample drawn from the population.

Inductive Inference



Hypothesis Tests Introduction

There are two *inferential* techniques a researcher uses to draw conclusions from a sample: *hypothesis testing* and the other is *interval estimation*.

Definition . . .

- " The *hypothesis test* or *test of significance* is based on the concept of proof by contradiction. The concept is composed of four parts.
 - " State the *null* and *alternative hypotheses*.
 - " Collect and assess the data.
 - " Calculate the *test statistic* and the probability of observing that particular value of the test statistic or one more extreme.
 - " *Make a decision* and state the conclusions.

Hypothesis Tests *Null & Alternative Hypotheses*

The purpose of a *test of significance* is to give a statement of the strength of the data against the null hypothesis.

Definitions . . .

- " The statement being tested (abbreviated H_0). Most often the null hypothesis is a statement of the status quo or no difference between two groups.
- " The statement the researcher alternative hypothesis (abbreviated H_a). The alternative hypothesis is also referred to as *research* or *working hypothesis*. The research hypothesis is usually collected. To verify the research hypothesis, an investigator tries to collect data in other words, a proof by contradiction.

Hypothesis Tests *Null & Alternative Hypotheses*

The *null hypothesis* can be thought of as . . .

- " a statement of no difference between two a statement of no difference between two populations or treatments;
- " a statement of no a statement of no difference between a statement of no difference between a sample or treatment and a hypothesized value;
- " a statement that the effect the researcher a statement that the effect is not present; and
- " a statement that the researcher hopes to find evidence against.

Hypothesis Tests *Null & Alternative Hypotheses*

Why is a *null hypothesis* a statement of no difference?

In most investigations a decision based on incomplete information is based on a sample not the entire population.

Because decisions are made based on incomplete information, a researcher can never prove a claim; rather, a researcher can only *disprove* a claim.

Hypothesis Tests *Null & Alternative Hypotheses*

To see the consequences of the hypothetical example.

Suppose Company X manufactures marbles using an automated process. The output is automatically collected in a single container. Millions of marbles are produced each day. Now suppose only white marbles are produced each day. Now suppose only white marbles. In order to verify that the marbles are only white, the operator samples the container by pulling out several marbles several times throughout the day. One day the operator samples the container and finds a black marble. Has he *proven* that all the marbles are white? No, that all the marbles are white. The operator is seeing only a small portion of the millions of marbles. His conclusion can *only* depend on what he sees. On the other hand, if the operator had found a black marble, this clearly *disproves* the claim that all the marbles are white.

Hypothesis Tests Null & Alternative Hypotheses

In the preceding example the null hypothesis is

H_0 : *the marbles are all white,*

a statement of the status quo.

The alternative hypothesis is

H_a : *the marbles are not all white.*

The alternative hypothesis is why the data are being collected, the operator suspects the status quo is incorrect.

By not finding any nonwhite marbles, the operator cannot prove the null hypothesis is true; the operator can only conclude that there is not enough evidence to reject his claim.

Hypothesis Tests Null & Alternative Hypotheses

The above example illustrates that the *null hypothesis can never be proven*, but rather, it is either rejected, but rather, it is not rejected based on the information gathered from an investigation.

If the null hypothesis is not rejected, If the null hypothesis is not *not proven* that it is true, but that it is true, but rather, that it is only concluded there is insufficient information to reject the null hypothesis.

This is why the null hypothesis is the researcher hopes to find evidence collected in an attempt to reject the null hypothesis.

Hypothesis Tests Null & Alternative Hypotheses

It is important to remember....

Failing to reject the null hypothesis does not mean that the null hypothesis is true. It only means that there is not enough evidence to reject it. A difference may exist, but due to an insufficient number of experiments, insufficiently homogeneous experimental techniques, or imprecise measurement techniques, this investigation may yield data that do not reject the null hypothesis.

Hypothesis Tests Null & Alternative Hypotheses

There are two types of hypothesis tests in the statistical literature as *directional* (*one-tailed*) and *two-tailed* (*two-sided*, *nondirectional*). The type of *alternative hypothesis* determines the location of the region of rejection for the H_0 .

In a *directional alternative*, the researcher provides information to suggest the direction of the information. For example,

$$H_0: \mu = 0 \text{ versus } H_a: \mu > 0.$$

The researcher believes the population mean is greater than 0.

Hypothesis Tests *Null & Alternative Hypotheses*

In a *nondirectional alternative*, the researcher h, the research information to suggest a direction. For example,


$$H_o: \quad = 0 \quad \text{versus} \quad H_a: \quad 0.$$

The researchThe researcher doeThe researcher does not population mean will deviate from 0.

Hypothesis Tests *Null & Alternative Hypotheses*

The *null* and *alternative hypotheses* form a dichotomy where only one of the hypotheses is true. If the null hypothesis is rejected, then the alternative hypothesis must be true.

Probability is used to quantify evidence against the null hypothesis. Probability tells the researcher what would happen if the researcher what was generated the data were *repeated again and again*....



This describes a *frequency*-based view of probability and is the foundation of classical statistics (classical statistical techniques such as the analysis of variance and the *t*-test).

Hypothesis Tests *Test Statistic*

The evidence against the null hypothesis is quantified by calculating a *test statistic* and determining and determining the probability (commonly referred to as the *p-value*) of observing that particular result or observing that p

Definition . . .

In general, the decision to reject or not reject the null hypothesis is based on a single value that summarizes the data, called the *test statistic*. The test statistic is a variable, associated with a test statistic, associated with a probability (commonly probability (commonly of observing that particular extreme.

Hypothesis Tests *Test Statistic*

Definitions . . .

- " A *random variable* assigns an event a number on the real line. For example, the number of games a team will win in a given season is a random variable.
- " The mechanism used to assign a number to a random variable is a *probability distribution*. A value of a random variable is a number in the interval $[0, 1]$ for a given event. The probability distribution is represented by a function. For example, the Student's *t*-distribution is a probability distribution function.

Hypothesis Tests *Test Statistic*

For example . . .

Suppose there is an urn with 5 balls, 3 of which are black and 2 of which are red. In other words, the population is $\{B_1, B_2, B_3, R_1, R_2\}$.

We are going to employ simple random sampling. We are going to select samples of size 3. There are 10 different samples that we can construct.

Let Q represent the number of red balls in the sample of 3 balls. Q is the random variable.

What are the possible values Q can take on?

Hypothesis Tests *Test Statistic*

<i>Population</i> {B1, B2, B3, R1, R2}			
<i>Event</i>	B1, B2, B3	B1, B2, R1	B1, R1, R2
selecting 3 balls using simple random sampling without replacement		B1, B2, R2	B2, R1, R2
		B1, B3, R1	B3, R1, R2
		B1, B3, R2	
		B2, B3, R1	
Set of possible outcomes		B2, B3, R2	
Values of the Random Variable	0	1	2
q = the # of red balls			
Probability of observing that value of the Random Variable	1/10	6/10	3/10

Hypothesis Tests *Test Statistic*

We have three pieces of information here,


1. a random variable,
2. all possible values of the random variable,
- 3.3. the probability of occurrence the probability of occurrence

These are the three pieces of information These are the
construct a *probability distribution* for a random
variable.

Probability Distribution for the Variable Q			
q	0	1	2
$P(Q = q)$	1/10	6/10	3/10

Hypothesis Tests *Test Statistic*

A histogram or
probability distribution function (PDF)
of the random variable on page 17.



Hypothesis Tests *Test Statistic*

Interpretation: of the PDF.

- " TheThe probability of observing 0 red balls when selecting a sample of 3 balls is

$$P(Q = 0) = 0.1.$$

- " TheThe probabilitThe probabilityThe probability of ob selecting a sample of 3 balls is

$$P(Q = 1) = 0.6.$$

The PDF is a *relative frequency histogram*

Definitions . . .

- " *Frequency* is the number is the number o is the number of observations in a particular particular class. The relative frequency is expressed as a preexpressed as a frequency.

Hypothesis Tests *Test Statistic*

There is an alternative way we can express the probability information about the random variable.

The *probability distribution* answers the question,

What is the probability that the variable *Q* equals a specific value?

Sometimes we are interested in a different question.

For example,

What is the probability that the variable *Q* is greater than or equal to a specific value?

Hypothesis Tests *Test Statistic*

Given the original *probability distribution*, the question,

What is the probability that the variable Q is equal to a specific value? , is easy to answer.

Probability Distribution for the Variable Q

q	0	1	2
$P(Q = q)$	$1/10$	$6/10$	$3/10$

Answer to the New Question for the Variable Q

q	0	1	2
$P(Q \leq q)$	$1/10$	$7/10$	$10/10$

The answer to this new question about the probabilities that precede the value of interest.

Hypothesis Tests *Test Statistic*

The table that gives the answers to the questions for the variable Q has a formal name.


It is called a *Cumulative Distribution Function* (CDF).

CDF for the Variable Q

q	0	1	2
$P(Q \leq q)$	1/10	7/10	10/10

Hypothesis Tests *Test Statistic*

The CDF is another way to think of a random variable. The CDF can be thought of as a *cumulative relative frequency* histogram.



Hypothesis Tests *Test Statistic*

Interpretation: of the CDF.

- " The probability of observing 1 or fewer when selecting a sample of size 3 is

$$P(Q \leq 1) = 0.7.$$

- " The probability of observing 2 or fewer when selecting a sample of size 3 is

$$1 - P(Q = 0) = 1 - 1/10 = 0.9.$$


The CDF is a *relative cumulative frequency histogram*.

Definitions . . .

- " The cumulative frequency is the number of observations less than or equal to a given value. The relative cumulative frequency is the cumulative frequency expressed as a proportion or percent of the total frequency.

Hypothesis Tests *Test Statistic*

What makes a *test statistic* different from other statistics?



A test statistic is a decision maker.

It has a probability distribution with it, whereas other statistics such as the mean and standard deviation do not. *Test statistics* are specific to the hypothesis being evaluated. For example, the commonly used test statistic, the *t*-test, evaluates whether a sample mean is different from the hypothesized mean or whether hypothesized means come from the same population.

Figure 1 on page 27, shows the population, sample statistic, and test statistic.

Figure 1. Distributions of Samples Means and t -Test Statistics

Hypothesis Tests *Test Statistic*

Figure 1 shows in order to make inferences, Figure 1 shows how we use probability in decision making,

- " the characteristic of interest is first summarized using a statistic, in this case the mean;
- " the statistic is then transformed where
- " the test statistic follows a known probability distribution.

The table on page 29 lists some test statistics and the null hypotheses they evaluate.

Hypothesis Tests *Test Statistic*

Statistical Test	Null Hypothesis	Test Statistic	Assumptions
Z-test	population mean $H_0: \mu = \mu_0$	$\frac{\bar{x} - \mu_0}{\frac{s}{\sqrt{n}}}$	" population variance, σ^2 , is known " population distribution is normal
Z-test	two population means $H_0: \mu_1 = \mu_2$	$\frac{\bar{x}_1 - \bar{x}_2}{\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}}$	" population variances, σ_1^2, σ_2^2 , are known " population distributions are normal
t-test	population mean $H_0: \mu = \mu_0$	$\frac{\bar{x} - \mu_0}{\frac{s}{\sqrt{n}}}$	" population variance, σ^2 , is <i>not</i> known " population distribution is normal
t-test	two population means $H_0: \mu_1 = \mu_2$	$\frac{\bar{x}_1 - \bar{x}_2}{\sqrt{\frac{s_p^2}{n_1} + \frac{s_p^2}{n_2}}}$	" population variances, σ_1^2, σ_2^2 , are <i>not</i> known " population distributions are normal " sample variances are equal,
t-test	paired comparison $H_0: \mu_d = 0$	$\frac{\bar{d} - \mu_d}{\frac{s_d}{\sqrt{n}}}$	" paired observation (2 observations on the same EU or an RBD with 2 EUs / block) " population distribution is normal
χ^2 -test	population variance $H_0: \sigma^2 = \sigma_0^2$	$\frac{(n-1)s^2}{\sigma_0^2}$	" population variance, σ^2 , is <i>not</i> known " population distribution is normal
F-test	two population variances $H_0: \sigma_1^2 = \sigma_2^2$	$\frac{s_1^2}{s_2^2}$	" population variances, σ_1^2, σ_2^2 , are <i>not</i> known " population distributions are normal

Hypothesis Tests *Test Statistic*

One statistic is the t -statistic,

$$\frac{\bar{x} - \mu_0}{\frac{s}{\sqrt{n}}} \quad (1)$$

It is a measure of the deviation of a sample mean from a hypothesized mean.

A theorem states that,

if all possible samples of size n are drawn from a normal population and if for each sample the calculated frequency values follows the *Student's t -distribution* with $(n - 1)$ degrees of freedom, where $(n - 1)$ is the number of degrees of freedom of s .

Hypothesis Tests *Test Statistic*

If the researcher is willing to make assumptions about the data, then the test provides the mechanism for the researcher to make inferences with *only one sample from the population*.

Using equation (1), the sample mean is assumed to follow a known probability distribution of the transformed variable. The PDF for the t -statistic is given by using the function

$$\frac{1}{\Gamma\left(\frac{k}{2}\right)} \frac{\Gamma\left(\frac{k}{2}\right)}{\sqrt{k\pi}} \left(\frac{k}{2}\right)^{\frac{k}{2}} \frac{1}{\left(1 + \frac{t^2}{k}\right)^{\frac{k+1}{2}}}$$

where $k = \text{degrees of freedom}$.

Hypothesis Tests *Test Statistic*

A histogram of 1,000 t -values
with 24 degrees of freedom (df).

Hypothesis Tests *Test Statistic*

It is the relative frequency, distinguishes one t -distribution from another.

The relative frequencies of the t -distribution various values of the df are provided in line of the t -table represents a particular value of the df .

Hypothesis Tests *Make a Decision*

The probability of observing a particular t -value is used to make a decision. But unfortunately a decision is made based on *incomplete information*, mistakes can be made. In hypothesis testing there are two potential decisions:

to reject or fail to reject the null hypothesis.

Each decision brings with it the possibility of an error being made.

- " A *type I error* is made if the null is rejected if the null is in fact true.
- " A *type II error* is made if the null is not rejected when it is in fact false.

The table of the next page illustrates the errors that can be made in hypothesis testing.

Hypothesis Tests *Make a Decision*

Potential Errors in Hypothesis Testing

Make the DECISION:	The NULL HYPOTHESIS is:	
	True	False
<i>Not to Reject</i> the Null Hypothesis	<i>Correct Decision</i> (1)	<i>Incorrect Decision</i> Type II Error ()
<i>to Reject</i> the Null Hypothesis	<i>Incorrect Decision</i> Type I Error ()	<i>Correct Decision</i> Power (1)

Hypothesis Tests *Make a Decision*

Definition . . .

The decision to reject the null hypothesis is in fact true is referred to as the type of decision error can be made only if the null hypothesis is in fact true.

If the null hypothesis is in fact true, the probability of making a type I error is denoted by symbol α . This is commonly referred to as *statistical significance level*.

Hypothesis Tests *Make a Decision*

Based on the probability of observing a particular value of the test statistic, the test is statistically significant or not statistically significant.

Definition . . .

This probability of observing a particular value of the test statistic is referred to as the *p-value*.

Traditionally, if the

$$p\text{-value} < \text{type I error (} \alpha \text{)},$$

the result is labeled as *statistically significant*, and the null hypothesis is rejected. A *statistically significant* result is one that is unlikely to occur by chance.

Hypothesis Tests *Make a Decision*

Definitions . . .

The decision not to reject the null hypothesis when it is in fact false, is referred to as a type I error. This type of decision error can occur if the null hypothesis is in fact false.

If the null hypothesis is in fact false, the probability of making a type II error is denoted by the Greek symbol β .

The power of a test, $1 - \beta$, is the probability of rejecting the null hypothesis when it is in fact false.

Hypothesis Tests *Make a Decision*

The relationship between α and β is illustrated in Figures 2 and 3.

Figure 2. Rejection / Acceptance Regions for the Hypotheses $H_0: \mu = 10$ versus $H_a: \mu > 10$ ($\alpha = 0.10$)

Hypothesis Tests *Make a Decision*

In this scenario, the null hypothesis is

$$H_0: \quad = 10.$$

The variability in the sample mean \bar{x} illustrates the fact that each time a sample is selected from the population a different value of the sample mean is calculated, and therefore a different value of the sample standard deviation is calculated.

The sample mean \bar{x} is based on n observations, therefore each time a sample is computed, the value of the mean will vary.

Hypothesis Tests *Make a Decision*

Figure 3. Rejection / Acceptance Regions for the Hypotheses $H_0: \mu = 10$ versus $H_a: \mu > 10$ ($\alpha = 0.01$)

Hypothesis Tests *Making a Decision*

In Figures 3 and 4:

- " The area under the *null distribution* (solid line) to the right of the vertical line is α .
- " The area to the right of the vertical line and enclosed by the null distribution is α .
- " The area under the *alternative distribution* (dashed line) to the left of the vertical line is β .
- " The area to the left of the vertical line and enclosed by the alternative distribution is β .

Hypothesis Tests *Making a Decision*

Going from Figures 2 to 3 notice that:

- " The probability of making Type I error (0.10) as the probability of making Type II error increases (0.008 to 0.0134).
- " This illustrates the *inverse relationship between type I and type II errors*. As long as the significance level remains constant, a decrease in type I error is accompanied by a simultaneous increase in type II error and vice versa.

Hypothesis Tests *Making a Decision*

Significant in the statistical sense does not mean *important*. It simply means *not* by chance. What is important is *practical significant difference*.

Definition . . .

The *practical significant difference* quantifies the evidence against the null hypothesis using a *p-value*.

A practically significant result is a difference between populations that influences the actions of the researcher.

Hypothesis Tests *Example Two-Sided*

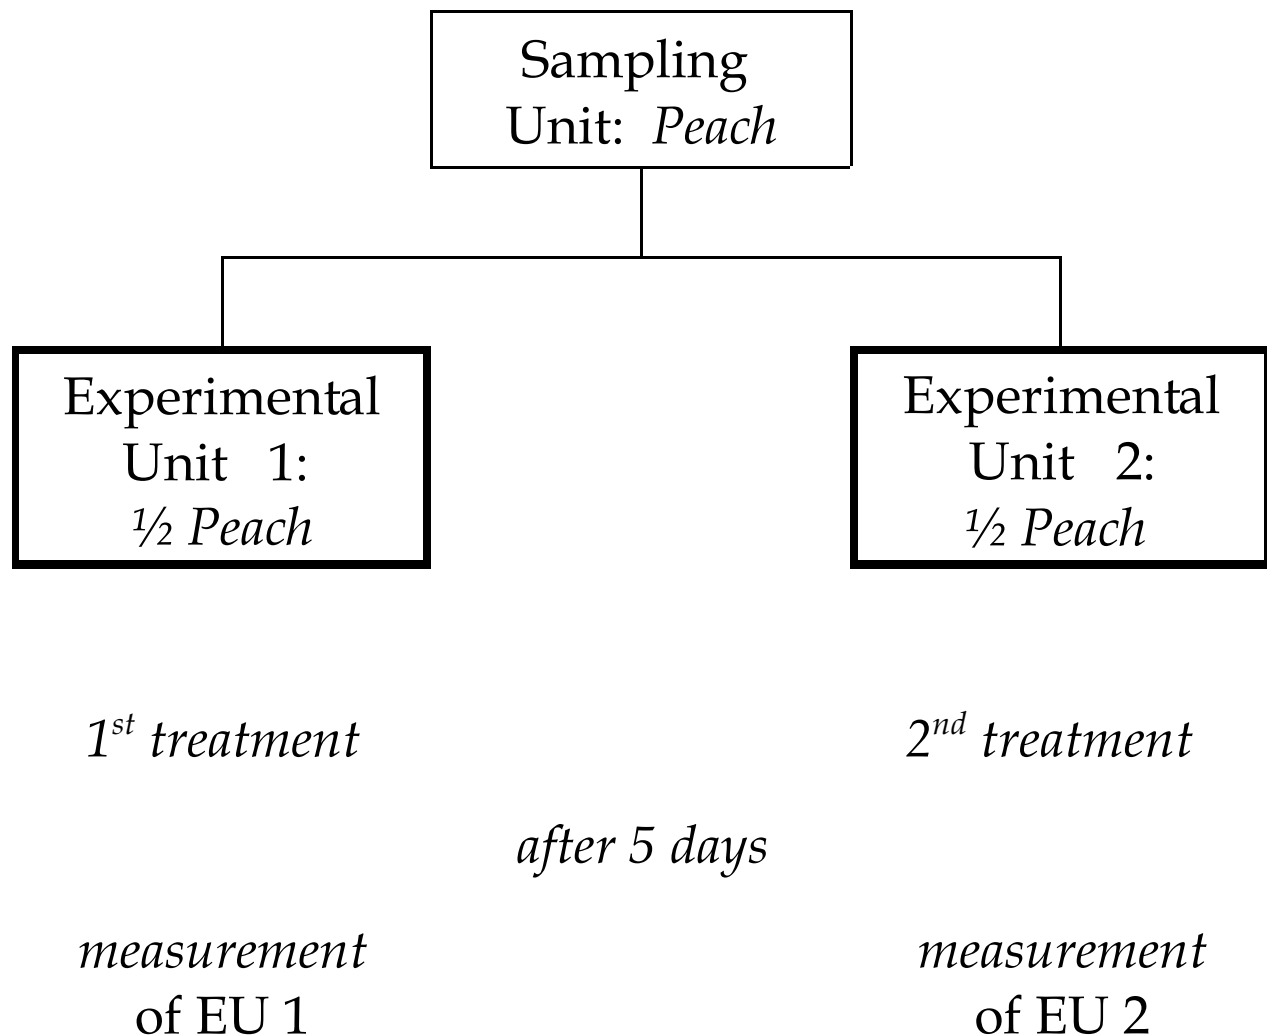
Let's see how this works with an example.

Say a researcher is interested in how the sugar content in peaches is affected by different storage conditions.

- " Peaches of the same variety and age are randomly selected from an orchard.
- " Each peach is cut in half and randomly assigned to one of two storage treatments.
- " After five days each peach has a sugar content.

Hypothesis Tests *Example Two-Sided*

Since each *sampling unit* receives
both treatments the data are paired.



Hypothesis Tests *Example Two-Sided*

Let,

x_{i1} = the peach half from the i^{th} peach receiving peach received first storage treatment, and

x_{i2} = the peach half from the peach half received second storage treatment.

One way to compare the relative effectiveness of two storage treatments is to compare the sugar content within a peach.

Let,

μ_D = the population mean difference between storage treatments for a particular peach.

Hypothesis Tests *Example Two-Sided*

Let,

D_i = represent the difference in concentration between i^{th} peach.

Assuming $n = 10$ sample units were taken symbolically we have,

$$D_1 = x_{11} - x_{12}$$

$$D_2 = x_{21} - x_{22}$$

•

•

•

$$D_n = x_{n1} - x_{n2}.$$

One way to compare the two treatments is to look at these differences and see whether the average difference is greater than or less than zero.

Hypothesis Tests *Example Two-Sided*

Our estimator for the population mean, μ_D , is,

$$\bar{D} = \frac{1}{n} \sum_{i=1}^n D_i,$$

the sample mean of our n differences. The difference data are used to calculate the statistic and this data is transformed into a known probability distribution.

Hypothesis Tests *Example Two-Sided*

Table 1. Sugar Concentration in Peaches.

<i>Peach</i>	<i>1st</i> <i>Treatment</i>	<i>2nd</i> <i>Treatment</i>	<i>Differences</i>
1	$x_{1,1}=23.69$	$x_{1,2}=30.25$	$D_1= 6.56$
2	$x_{2,1}=42.61$	$x_{2,2}=26.07$	$D_2= 16.54$
3	$x_{3,1}=30.73$	$x_{3,2}=39.25$	$D_3= 8.52$
4	$x_{4,1}=26.73$	$x_{4,2}=19.26$	$D_4= 7.47$
5	$x_{5,1}=34.84$	$x_{5,2}=35.22$	$D_5= 0.38$
6	$x_{6,1}=38.72$	$x_{6,2}=34.14$	$D_6= 4.58$
7	$x_{7,1}=45.29$	$x_{7,2}=33.43$	$D_7= 11.86$
8	$x_{8,1}=26.87$	$x_{8,2}=30.36$	$D_8= 3.49$
9	$x_{9,1}=26.24$	$x_{9,2}=22.13$	$D_9= 4.11$
10	$x_{10,1}=37.64$	$x_{10,2}=19.37$	$D_{10}= 18.27$

Hypothesis Tests *Example Two-Sided*

Since there is no information to suggest one will perform better than the other, this is a *two-sided hypothesis test*. The hypotheses are:

$$H_0: \mu_D = \mu_{D0} \text{ versus } H_a: \mu_D \neq \mu_{D0}$$

where,

μ_D = the population mean of differences between particular peach, and

μ_{D0} = the hypothesized population mean under the null hypothesis, which for this scenario is 0.

Hypothesis Tests *Example Two-Sided*

Things Things to think about Things to think about are how much the H_0 will we insist on?

That That is, what That is, what is the *significance level* committing a *type I error*?

For an $\alpha = 0.10$,

we are requiring that the data give evidence against H_0 so strong that it would happen no more than 10% of the time (2 times in 10) if the fact true.

For an $\alpha = 0.01$,

we are insisting on stronger evidence against H_0 , evidence so strong that it would happen no more than 1% of the time (1 time in 100) if the fact true.

Hypothesis Tests *Example Two-Sided*

Since we are testing a *population* where,

- " the *variance is unknown*, and
- " the sample size is *small* (< 30),

the *t-test* statistic is appropriate,

$$\sqrt{\frac{\bar{x} - \mu_0}{s^2/n}} .$$

Hypothesis Tests *Example Two-Sided*

Hypothesis Tests *Example Two-Sided*

Most classical procedures are based on the assumptions of the random variable, of the validity of the analyses depends on the assumptions. For the paired t -test the assumptions are:

- " D_i s are the result of paired measurements,
- " D_i s are independent, and
- " D_i s are normally distributed,
- " where the assumption of where the assumption of are no *outliers* .

Hypothesis Tests *Example Two-Sided*

The *graphical methods* of EDA of EDA provide EDA diagnostic tools for confirming assumptions are not met assumptions are not met actions.

Graphical displays meet the need to see the behavior of the data, to reveal the unexpected features, such as *outliers*; and confirm or disprove *assumptions*, such as the distributional assumptions of normality.

Hypothesis Tests *Example Two-Sided*

Two types of plots commonly utilized in EDA are, histograms and boxplots. From these histograms and sample differences, the data appear to be approximately normally distributed.

For comparison, consider histograms of random samples ($n=10$) from the standard normal distribution and random samples ($n=10$) from the standard normal distribution.

Hypothesis Tests *Example Two-Sided*

This illustrates that random samples from This distribution can look very *abnormal* when the sample size is *small*.

Hypothesis Tests *Example Two-Sided*

The boxplot for the sample data provide evidence against the assumption the data are normally distributed.

For comparison, let's look at the boxplot used to construct the histograms on page 56.

Hypothesis Tests *Example Two-Sided*

Again, Again, Again, random Again, random samples from a normal distribution appear somewhat *non-normal* when the sample size is *small*.

Hypothesis Tests *Example Two-Sided*

Why should we be concerned about outliers ?

By looking at the test statistic,

$$\frac{\bar{x} - \mu_0}{\frac{s}{\sqrt{n}}},$$

we can see that the behavior of the test statistic is driven by the sample mean and sample standard deviation.

Since they are not resistant measures of location and shape. The sample mean and sample standard deviation are both impacted by the presence of outliers.

Hypothesis Tests *Example Two-Sided*

What are *outliers* . . .

- " An *outlier* in any graph of data is an individual observation that falls outside the graph.

(David Moore, *Statistics Concepts and Controversies*)

- " Highly suspect observations in a sample.

(B. J. Winer, *Statistical Principles in Experimental Design*)

- " . . . observed values which . . . observed values remote from the main body of observations be discarded as being erroneous miscalculations, etc.

(Rupert Miller, *Beyond ANOVA, Basics of Applied Statistics*)

Hypothesis Tests *Example Two-Sided*

The impact of \bar{x} outliers on \bar{x} and s_D depends on the sample size and the variability.

If the sample size is large enough, then the impact of outliers on the value of the t -statistic toward the outlier is small.

In this case,

one is more likely to reject the H_0 due to the inflated value of the test statistic.


Hypothesis Tests *Example Two-Sided*

If the sample size and the variance are large, then outliers can actually have the impact of pulling the t -statistic toward the hypothesized value.

In this case,

one is *less likely to reject* the H_0 due to the deflated value of the test statistic.

Recommendation



Even when there is no assignable cause for the outlier and it is a true outlier, it does not necessarily mean the data should be changed or omitted. Prand should be all analyses with and without the outlier(s).

When there are no differences between the two sets of analyses the dilemma is resolved.

When there is a difference, an interpretation of both of both results should be provided in conclusions.

Hypothesis Tests *Example Two-Sided*

Based on the histogram (page 57) and boxplot (page 59) there are concerns about the assumption of normality, therefore we cannot use the t -test.

Note:

There are many *qualitative* (EDA) and *quantitative* (CDA) statistical techniques that do not require distributional assumptions and are robust to outliers. This is the topic of Workshop 10: Exploratory and Confirmatory Data Analysis.

Hypothesis Tests *Example Two-Sided*

The D_i values in Table 1 (page 50) are summarized below, below, this information is used below, this information is used

\bar{D}	s_D	D_0	n
4.39	9.29	0	10

The observed value of the test statistic is

$$t_{\text{obs}} = \frac{\bar{D} - D_0}{s_D / \sqrt{n}} = 1.49.$$

The next step is to find the *critical values* for a two-sided test based on the t -distribution with $\alpha = 0.10$, $\pm t_{(9, 0.10/2)}$.

Hypothesis Tests *Example Two-Sided*

The *critical values* partition the quantiles partition into two regions, the region where the H_0 is not rejected.

Since $t_{\text{obs}} = 1.49$ is less than the absolute value of the critical values $|t_{(9, (0.10/2)}| = 1.83$, we fail to reject the hypothesis.

Hypothesis Tests *Example Two-Sided*

The *critical values* can be found using the quantiles of the t -distribution and z values.

Example of a t -Table

df	=0.1*	=0.05	=0.025	=0.01	=0.005	=0.001
...
6	1.440	1.943	2.447	3.143	3.707	5.208
7	1.405	1.895	2.365	2.998	3.499	4.785
8	1.397	1.860	2.306	2.896	3.355	4.501
9	1.383	1.833	2.262	2.821	3.250	4.297
10	1.372	1.812	2.228	2.764	3.169	4.144
...

*The area under the curve to the right of the t -value (quantile).

t -tables are usually found in statistical methods books.

Hypothesis Tests *Example Two-Sided*

The hypotheses can be thought of as the probability of observing a value of t_{obs} or one more extreme.

- " A p -value can be thought of as an value can be thought of as a significance level.
- " Calculating a p -value gives more information than simply comparing a test statistic to a *critical value(s)*.

The p -value is the probability of observing a value of t_{obs} or one more extreme.

Hypothesis Tests *Example Two-Sided*

For the peach investigation, the p -value is the probability of observing a test statistic as extreme as the one calculated, 4.39, or less than -4.39 when the true population difference is 0.

Since 4.39 corresponds to a t -value, we need to calculate the probability of observing this result or one larger, given the true value.

Most often the p -value is calculated using a statistical package that calculates the probability distribution.

Hypothesis Tests *Example Two-Sided*

For a $t_{\text{obs}} = 1.49$, the area under the curve to the right of t_{obs} is 0.085. For a two-sided test, the value is 0.170.

Hypothesis Tests *Example Two-Sided*

Since our alternative hypothesis is *two-sided*, we now consider results that are equally extreme respect to the other side of the null hypothesis.

Stated another way, this means we stated sample means that correspond to a $t_{obs} = 1.49$.

Hypothesis Tests *Example Two-Sided*

The probability of observing this set of data more extreme with respect to the null hypothesis is given by,

the observed p-value =

$$\text{Prob}(t_9 \leq -t_{obs}) + \text{Prob}(t_9 \geq t_{obs}) =$$

$$\text{Prob}(t_9 \leq -1.49) + \text{Prob}(t_9 \geq 1.49) =$$

$$0.17$$

Since 0.17 is greater than our significance level, we conclude there is not enough evidence to conclude the null hypothesis is false.

Hypothesis Tests *Conclusions*

Both approaches lead to the same conclusion. The only difference is:

- " performing the test using a *critical value* versus whether or not we reject H_0 whereas
- " the *p-value* quantifies the weight of evidence a particular sample provides against the H_0 .

Hypothesis Tests *Example One-Sided*

Now let's change the setup a bit.

Suppose we want to test a new storage against a standard one currently in a store. The new is not willing to employ the new is not willing to employ performs as well as or better than the current one.

This information is used to This information is used to This in a *one-tailed hypothesis*. . If we let storage treatment 2 be the standard treatment, where

$$D_i = x_{i1} - x_{i2},$$

then hypotheses are not,

$$H_o: \mu_D \leq \mu_{D_o} \text{ versus } H_a: \mu_D > \mu_{D_o}.$$

Hypothesis Tests *Example One-Sided*

For a less than 10% of hypothesis when it is in fact value for this *one-sided* test is $t_{(9, 1 - 0.10)} = 1.38$.

The difference between the *one-sided* test and the *two-sided* test is that with the *one-sided* test only one direction are considered important.

Hypothesis Tests *Example One-Sided*

In a one-sided test the entire significance level is on one side of the distribution.

It is easier to see a statistically significant difference in a *one-sided* test as opposed to a *two-sided* test since we are only looking in one direction of the alternative hypothesis.

Hypothesis Tests *Example One-Sided*

For the one-sided test, the observed statistic is still 1.49. Therefore, the test statistic is still 1.49. The test statistic is compared to the critical value at the 0.10 significance level since $t_{obs} > t_{(9, 1 - 0.10)} = 1.38 = 1.3$ where 1.38 is the critical value.

Hypothesis Tests *Example One-Sided*

What is the observed p -value that corresponds to -1.49 ?

The observed p -value = 0.085. Since the observed p -value is smaller than the *significance level*, conclude there is sufficient evidence at the 0.10 level.

Hypothesis Tests *Example One-Sided*

All other things being equal, a *one-sided test* is *one-sided* *more powerful* than a *two-sided test*, when the researcher has information to specify the direction of the effect.

By *more powerful*, it is meant that the probability of rejecting the null when it is in fact false is greater.

Interval Estimation *Introduction*

The *confidence interval* is the other *inferential* technique a researcher can use to draw conclusions from sample data.

Definitions . . .

- " The *confidence interval* includes a point estimate of the population parameter (for example, the population mean) accompanied by a measure of the associated uncertainty (the margin of error).
- " The two extreme points in a confidence interval are the lower and upper *confidence bounds*. The *confidence level* is the range of values within which there is a specified level of confidence that the true population parameter will fall.

Interval Estimation *Introduction*

Like the alternative hypothesis, there are two types of confidence intervals, referred to in the literature as *one-tailed* (*one-sided* or *directional*) and *two-tailed* (*two-sided* or *nondirectional*).

A *two-sided confidence interval* for a population parameter consists of the sample estimate for the parameter plus and minus the *margin of error*,

$$\{\text{sample estimate} \pm \text{margin of error}\}$$

A *one-sided confidence limit* for a population parameter consists of the sample estimate minus the *margin of error* or the sample estimate plus the *margin of error*.

$$\begin{aligned} & \text{lower confidence limit (LCL)} \\ &= \{\text{sample estimate} - \text{margin of error}\}, \text{ or} \end{aligned}$$

$$\begin{aligned} & \text{upper confidence limit (UCL)} \\ &= \{\text{sample estimate} + \text{margin of error}\}. \end{aligned}$$

Interval Estimation *Introduction*

Definition . . .


- " TheThe measure of confidenceThe measure of confidence intervalinterval includes theinterval includes the true population parameter. This is called the confidence coefficient. The confidence coefficient is the proportion of all possible sample confidence intervals that contain the true population parameter. The confidence coefficient is represented symbolically by C .

Interval Estimation *Introduction*

Interval estimation and hypothesis testing are related, in fact, a confidence interval can be used to conduct a hypothesis test.

For example, to test $H_0: \mu = \mu_0$ versus $H_a: \mu \neq \mu_0$ using interval estimation, the confidence interval is calculated from the sample data and

if the interval contains μ_0 , the null hypothesis is accepted; otherwise, it is rejected.



the null hypothesis is accepted; otherwise, it is rejected.

A $(1 - \alpha)100\%$ two-sided confidence interval for the population mean is equivalent to a two-sided hypothesis test at the α level of significance. This relationship is illustrated in Figure 4 using the hypotheses $H_0: \mu = \mu_0$ versus $H_a: \mu \neq \mu_0$.

Interval Estimation *Introduction*

Figure 4. Hypothesis Testing *versus* Interval Estimation: Rejection / Acceptance Regions for the Hypotheses $H_0: \mu = 0$ *versus* $H_a: \mu \neq 0$, ($\alpha = 0.05$)

Interval Estimation *Introduction*

For this hypothetical illustration, ten were randomly selected from ten were randomly selected with mean 0 and variance 2; with mean 0 and variance 2. For these samples, the t -test statistic used to conduct the hypothesis test and the 95% confidence interval were calculated.

The null hypothesis is not rejected when

- " the line segment crosses the solid line segment crosses the hypothesized value zero, or
- " when the filled circle is within the fail-to-reject region bounded by the solid vertical lines at ± 2.23 .

Interval Estimation *Introduction*

For each of the 50 samples,

" *the conclusions drawn based on either the conclusions drawn from the test or the confidence interval are identical.*

In only two cases did the sample data produce:

aa value of the test statistic a value of the test statistic outside the critical region; and

aa confidence interval that did not contain the hypothesized mean.

Interval Estimation *Introduction*

A common *misinterpretation* of the confidence interval is that . . .

a 90% *confidence interval* implies there is a 10% chance that the true parameter lies within the interval.

For a $(1 - 0.10)$ confidence coefficient, the probability that the interval will contain the unknown true parameter is 0.90.


The probability, 0.90, is a measure of the probability that the sample is drawn, that the true population parameter is estimated to be either in or not in the interval.

Interval Estimation *Introduction*

An analogy can be made to flipping a coin.

Before a fair coin is flipped, there is a probability of 0.5 that the coin will display heads. This probability is associated with the *random event* of flipping the coin. Once the random event has occurred, there is no longer any uncertainty of confidence that the coin will display heads.

Stated another way, if samples are repeatedly drawn from the same population, the confidence interval will contain the true unknown population parameter with a relative frequency of those intervals that would approach 95 percent.



This is a *frequency*-based view of probability

Interval Estimation *Introduction*

This is seen in the hypothetical Figure 4.

The hypothesized mean of zero (population parameter) is contained in 48 out of the 50 confidence intervals, or 96% of the time, the interval contains the population parameter.

Interval Estimation *Example Two-Sided*

The null hypothesis of the peach investigation be evaluated by constructing a *confidence interval*.

If we let α represent the *significance level*, then the $100(1 - \alpha)\%$ *confidence interval* will yield results.

Since $\alpha = 0.10$, a $100(1 - 0.10) = 90\%$ *confidence interval* using the 10 sample differences is constructed.

Interval Estimation *Example Two-Sided*

A 90% confidence interval for the

$$\bar{x} \pm t_{(n-1, (1-\alpha)/2)} \left(\frac{s}{\sqrt{n}} \right),$$

where,

- " \bar{x} is the parameter estimate,
- " $t_{(9, (1 - 0.10/2))}$ is the tabled value for a *2-sided confidence interval* using a *t*-distribution with $(10 - 1) = 9$ degrees of freedom and a confidence level of $(1 - 0.10) = 0.90$; and
- " $\left(\frac{s}{\sqrt{n}} \right)$ is the standard deviation of the sample differences divided by the square root of the sample size, the is referred to as the standard error of the mean differences, $\frac{s}{\sqrt{n}}$.

Interval Estimation *Example Two-Sided*

Why use $t_{(9, (1 - 0.10/2))}$?

- " For a 0.90 *confidence coefficient*, 10% of the area, 10% outside of the *confidence interval*.
- " Since this is a two-sided *t*-distribution is symmetric, the UCL and 5% of the area is below the LCL.
- " $t_{(9, (1 - 0.10/2))}$ is used rather than $t_{(9, 0.10/2)}$, in order to have a positive *t*-value. Note that,

$$t_{(9, (1 - 0.10/2))} = |t_{(9, 0.10/2)}|,$$

since the *t*-distribution is symmetric.

Interval Estimation *Example Two-Sided*

The t -value that has 5 percent of the value that has 5 percent of the area under the t -distribution curve to its right is $t_{(9, (1 - 0.10/2))} = 1.83$.

Interval Estimation *Example Two-Sided*

For the peach investigation (using the statistics on page 66),

$$\begin{aligned}
 &= t_{(9, (1 - 0.10/2))} \left(\frac{\bar{D}}{\sqrt{s_D}} \right) \\
 &= \left(\frac{\bar{D}}{\sqrt{s_D}} \right) \\
 &= 4.39 \pm 5.38 \\
 &= [-1.00, 9.77].
 \end{aligned}$$

The 90% *confidence interval* for the mean difference for the peach investigation is $[-1.00, 9.77]$.

Since the hypothesized mean, $\mu_{D_0} = 0$, falls within the 90% *confidence interval*, there is no evidence to reject the H_0 at the 0.10 *significance level*.

Interval Estimation *Conclusion*

For evaluating hypotheses of the form,

$$H_o: \mu = \mu_o \text{ versus } H_a: \mu \neq \mu_o$$

the $100(1 - \alpha)\%$ confidence interval is equivalent to a *two-sided* hypothesis test of the mean at the α level of significance.

Interval Estimation *Example One-Sided*

We can calculate a *confidence limit* for directional hypotheses,

$$H_o: \mu \leq \mu_0 \text{ versus } H_a: \mu > \mu_0.$$

For a *one-sided alternative* in the right distribution, we need to calculate the *confidence limit*,

$$LCL = \{\text{sample estimate} - \text{margin of error}\}.$$

If the hypothesized value, μ_0 , is less than the LCL then conclude there is a significant difference from the null hypothesis.

Interval Estimation *Example One-Sided*

For the peach investigation (using the statistics on page 66),

$$\begin{aligned}
 LCL &= \bar{x} - t_{(9, (1 - 0.10))} \left(\frac{s}{\sqrt{n}} \right) \\
 &= 4.39 - 1.383 \left(\frac{0.406}{\sqrt{10}} \right) \\
 &= 4.39 - 0.176 \\
 &= 4.214
 \end{aligned}$$

The 90% lower confidence limit for the mean diameter is 4.214.

Since the hypothesized mean $\mu = 4.06$, and $LCL = 4.214$, there is enough evidence to reject H_0 at the 0.10 significance level.

Interval Estimation *Example One-Sided*

Why use $t_{(9,(1-0.10))}$?

- " FForFor a 0.90 *confidence coefficient*, 10% of the area, 10% outside of the *confidence interval*.

- " SinceSince this is a *one-sidedone-sided limit* and the alternat inin the right tail, all 10% oin the right tail, all 10% of th LCL.

- " SinceSince the *t-dis-distribution-distribution -distribution coefficient* is used rather than th is used rather than the is use order that the *t-value* is positive.

Interval Estimation *Conclusion*

For evaluating hypotheses of the form,

$$H_0: \mu \leq \mu_0 \text{ versus } H_a: \mu > \mu_0$$

the 100(1 - α)% *lower confidence limit* for the mean μ is
equivalent to a *one-sided* hypothesis test at the α level of significance.

Interval Estimation & Hypothesis Tests *References*

- (1)(1) Moore, David S., (1997). *Statistics Concepts and Statistics Concepts and Controversies*, 4th edition. W. H. Freeman and Company.

A very well written book on statistical concepts that helps you understand data in the face of uncertainty, explore the data, and how to draw conclusions from data. There are a few equations, but rather the focus is on the thought process behind the arguments and how to be a critical consumer of scientific studies and surveys.

- (2) Utts, Jessica M., (1996). *Seeing Through Statistics*, Belmont CA: Duxbury Press.

Another very well written book on statistical concepts to help the reader understand scientific studies and surveys, to help the reader obtain accurate information from useless and misleading; and to interpret accurate information using inferential methods, by focusing on how to understand them rather than how to compute them. Utts uses a case-study approach to compute them. Utts uses a case-study interpretation of data that appear in the media and scientific literature.

- (3) Kanji, Gopal K. (1997). *100 Statistical Tests*, SAGE Publications, London.

This is a great reference that begins with aThis is a great reference that
The remainder of the book presents the vaThe remainder of the book presents the va
parametricparametric and non-parametric. Each statistical test iparametric and non-parametric
test,test, the limitations, how to conduct thetest, the limitations, how to conduct the ttest, the
example. This is a very useful reference.

- (4) Ott, Ott, Lyman (1988). *An Introduction to Statistical Methods*, 3rd edition, PWS-Kent Publishing Company, Boston.

This book is appropriate for someone with little or no exposure to statistical methods and data analysis. It assumes no prior knowledge of statistics.